

This Page Is Inserted by IFW Operations
and is not a part of the Official Record

BEST AVAILABLE IMAGES

Defective images within this document are accurate representations of the original documents submitted by the applicant.

Defects in the images may include (but are not limited to):

- BLACK BORDERS
- TEXT CUT OFF AT TOP, BOTTOM OR SIDES
- FADED TEXT
- ILLEGIBLE TEXT
- SKEWED/SLANTED IMAGES
- COLORED PHOTOS
- BLACK OR VERY BLACK AND WHITE DARK PHOTOS
- GRAY SCALE DOCUMENTS

IMAGES ARE BEST AVAILABLE COPY.

**As rescanning documents *will not* correct images,
please do not report the images to the
Image Problem Mailbox.**

(12) UK Patent Application (19) GB (11) 2 343 025 (13) A

(43) Date of A Publication 26.04.2000

(21) Application No 9823187.1

(22) Date of Filing 23.10.1998

(71) Applicant(s)

**International Business Machines Corporation
(Incorporated in USA - New York)
Armonk, New York 10504, United States of America**

(72) Inventor(s)

Paolo Squartini

(74) Agent and/or Address for Service

**G M Zerbi
IBM United Kingdom Limited, Intellectual Property
Department, Mail Point 110, Hursley Park,
WINCHESTER, Hampshire, SO21 2JN,
United Kingdom**

(51) INT CL⁷

G06F 1/00

(52) UK CL (Edition R)

G4A AAP

(56) Documents Cited

None

(58) Field of Search

UK CL (Edition Q) G4A AAP

INT CL⁶ G06F 1/00

Online: EPODOC

(54) Abstract Title

License management system

(57) A method and system for providing flexibility to a license management system. A license management system permits the concurrent use of n copies of a software program over a network comprising a plurality of client workstations, each client workstation having a copy of the software program installed thereon requiring an authorization from one of a plurality of S license servers each time the software program is used. For security reasons, the license management system requires that at least the integer majority of the plurality of license servers is active at any time. The method and system of the present invention allow to change the number of license servers, but impose the following limit: the difference between the maximum allowed number of servers belonging to the cluster ($S+d$) and the minimum allowed number of servers belonging to the cluster ($S-d$) must be less than the majority of the maximum number of server belonging to the cluster.




FIG. 1

At least one drawing originally filed was informal and the print reproduced here is taken from a later filed formal copy.

This print takes account of replacement documents submitted after the date of filing to enable the application to comply with the formal requirements of the Patents Rules 1995


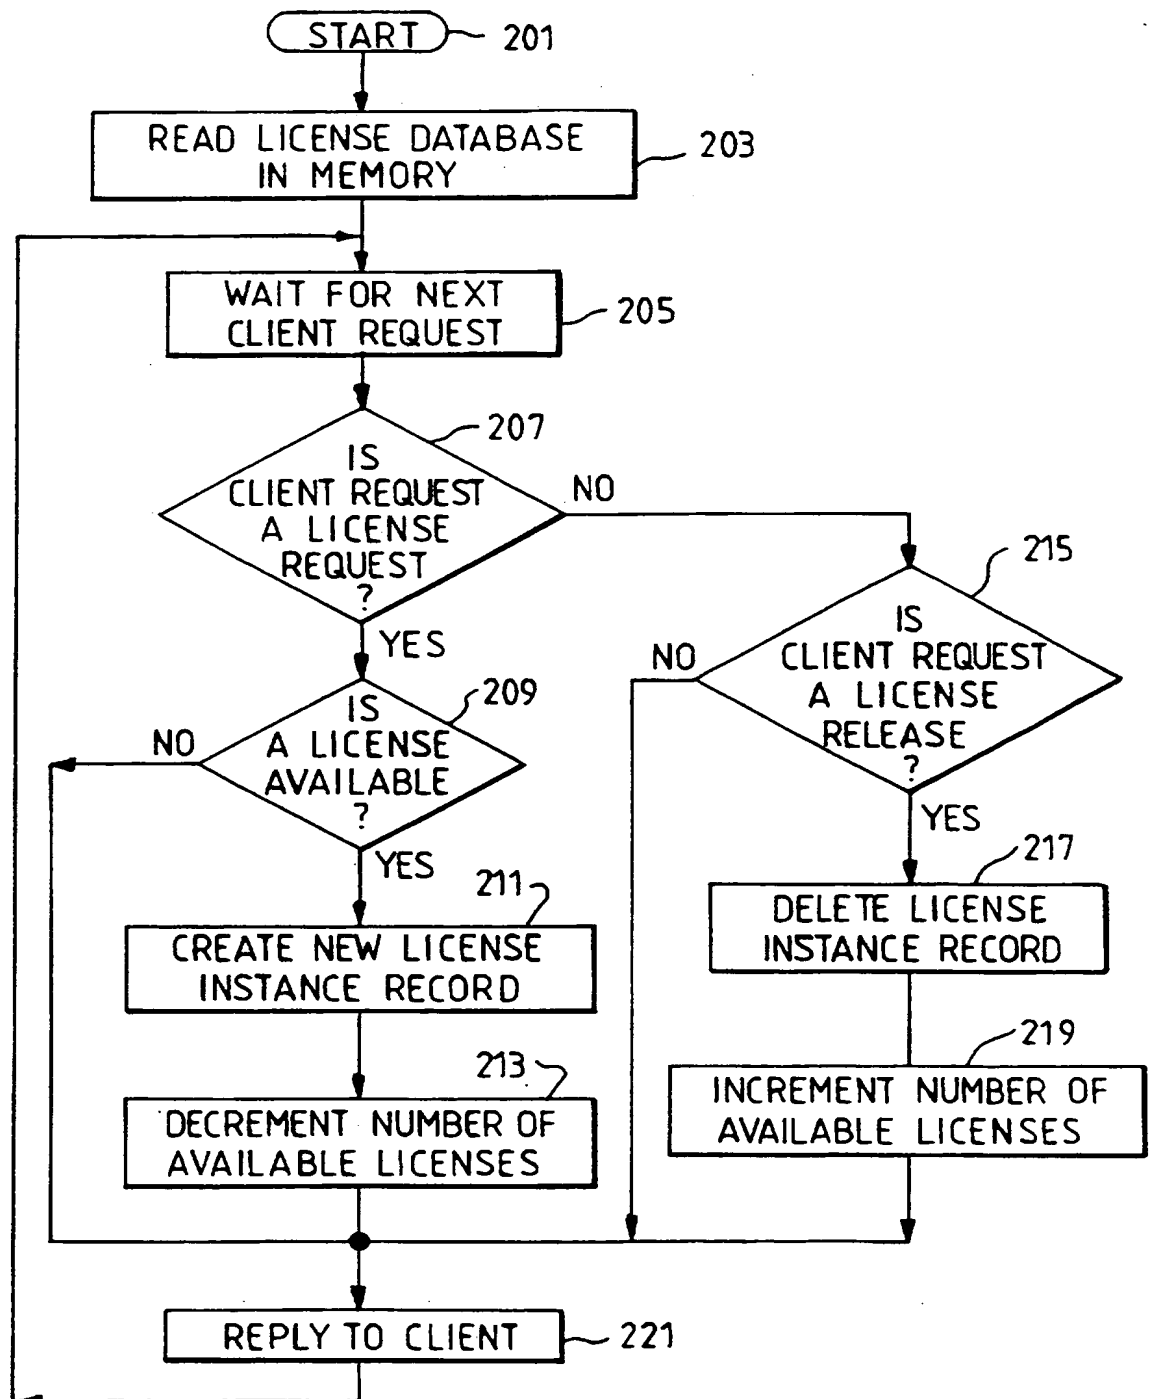
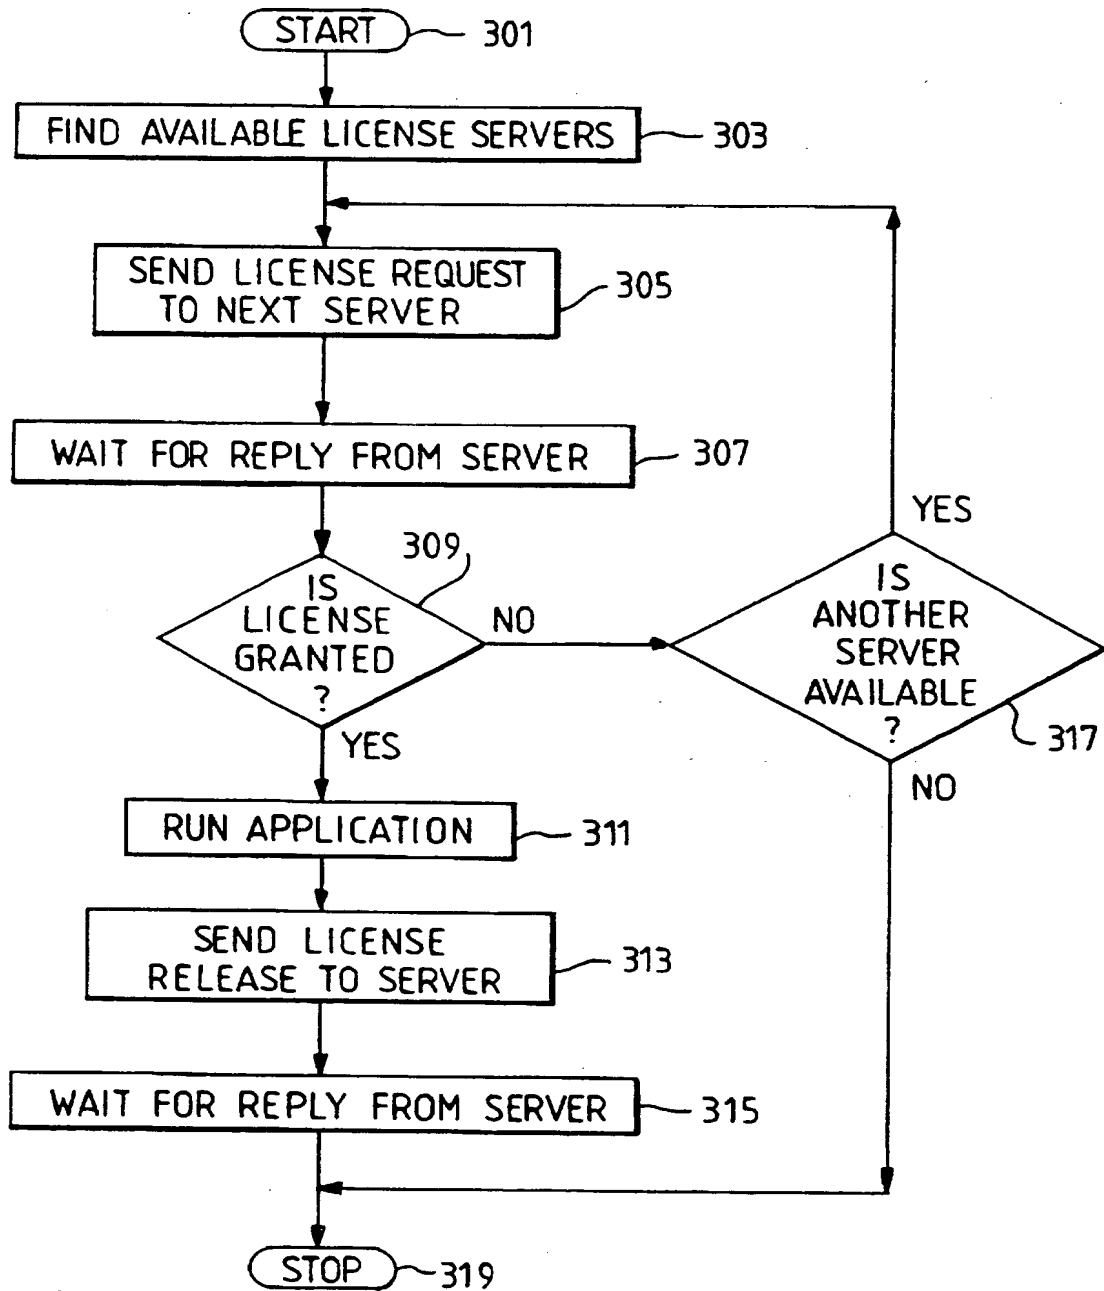
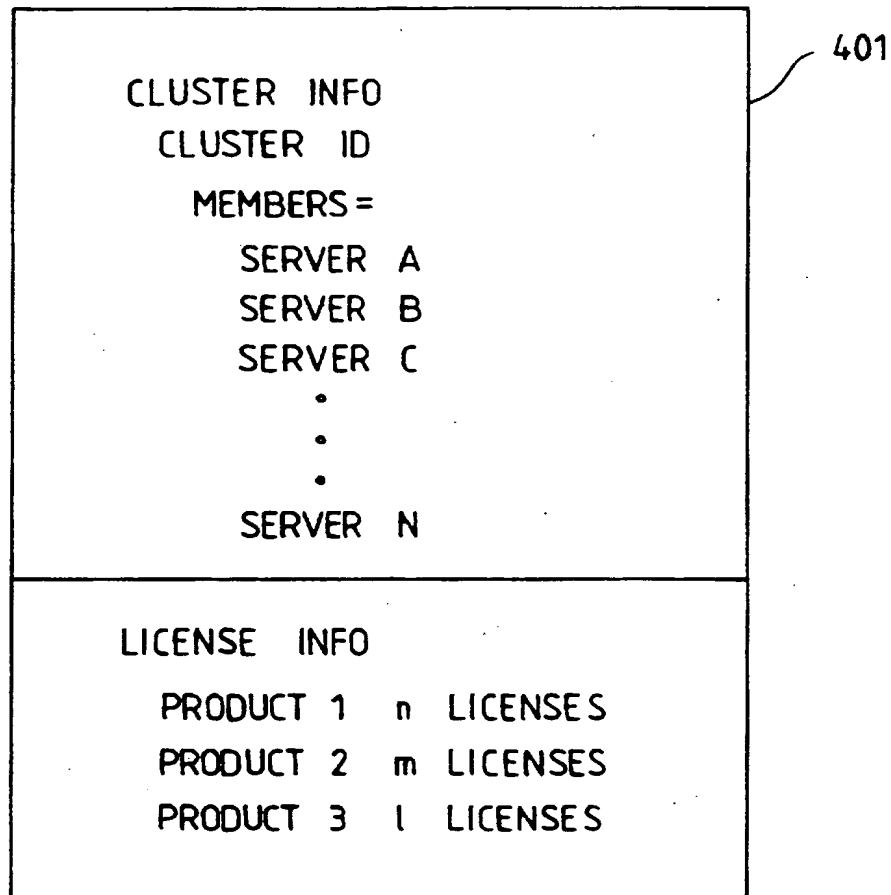


FIG. 1

FIG. 2

FIG. 3

EXAMPLE OF LICENSE DATABASE

FIG. 4




FIG. 5A





FIG. 5B



LAN SPLIT (2 LANS INSTEAD OF 1)
 1 CLUSTER WITH 2 SERVERS OUT
 OF 3. 1 PRODUCT n LICENSES
 1 CLUSTER WITH 3 SERVERS OUT
 OF 5. 1 PRODUCT n LICENSES
 TOTAL LICENSES = 2n

FIG. 5C

LICENSE MANAGEMENT SYSTEM

The present invention relates to license management systems and particularly to a method and system for providing flexibility to a license management system.

The licensing of computer software was traditionally accomplished by providing a copy of the software and a license for each computer which was authorized to use the software. The software could be generally used only on that computer, unless it was deleted from that computer and transferred to another one together with the license. With the advent of wide spread computer networks a more efficient solution was required. A known license management system allows a user to install a copy of a software program on N nodes of a network, but acquire only a limited number n licenses, where at any time, only the maximum number n copies of that program can be simultaneously run on the network. When all the available licenses are allocated, additional users requesting the use of that software must wait for a license to become available. This kind of license management system has a number of advantages for both the software vendor and the user, because it allows the user to purchase all and only the licenses really needed and, on the other hand, allows the vendor to fight software piracy.

An example of a state of the art license management system available on the market, is the License Use Management product of International Business Machines Corp.

In a typical network of interconnected computers with a license management system, as illustrated in Fig. 1, one or more of the nodes 101 act as license servers, while a plurality of nodes 103 act as clients of the license servers. The service provided by a license server 101 to its client 103 is that of granting or denying permission to run a given software program, according to the availability of a license record in the license server data base, and to the terms and conditions encoded in the license record itself. The license server usually creates and stores license records in the license data base upon processing license certificate files, which are provided by the software vendor and complement the software program to which they are related. This license data base must be locked in some way to the specific instance of the license server (hardware + software) to prevent malicious users from copying the license data base to another license server machine and multiplying by two the number of licenses for all the software products contained in the license data base. License certificate files may contain some encryption or checksum information that allow the license server to verify their authenticity and integrity.

The fact that a license management system is monitoring the use of a given software program should be as transparent as possible to the users of

that software program whereas it should be evident and beneficial to the administrator of licenses for that and other software programs. This consideration places a strong requirement on the license management system in terms of reliability and performance. The ideal license management system should be one which never causes software program failures because of its outage nor becomes a bottleneck for the software programs that it monitors.

In a license management system, "availability" is a measure of the degree to which the system can process and satisfy incoming requests (either granting or denying permission to run) within the time limits set by the served environment. High availability systems attempt to provide a continuous service within a particular operational window by minimising causes of failure and minimising recovery time when failures occur. Usually this requires a large degree of redundancy in system components, so that the continued operation of the entire system is protected from failure of any single component. The ultimate objective is to eliminate all single points of failure in the system. This can be accomplished by having redundant components or systems, and "availability management technology" that can automate the transfer of services to those redundant components when a failure occurs. Availability is a crucial feature of license management systems, since an outage of one or more license servers of a license management system can prevent many users from running their critical applications, due to a failure to acquire a license. An obvious solution to ensure good availability would be to use well known clustering techniques. In the network data processing field, a cluster is a set of independent processors (nodes), connected over the network. A cluster constitutes a sort of "black box" which provides certain services to end users. Like any ideal black box system, the end users do not need to know which node in the cluster they are connecting to. However, common clustering techniques, aimed at increasing the overall availability of the system through server redundancy, cannot be applied in a straightforward way to license management systems because of the secure nature of the license serving environment. A redundant license server cannot simply take over the amount of licenses, served by another failing server; it must also ensure that, in no circumstances, the total number of licenses concurrently served can exceed the total number of available licenses, stored into the license authorization record.

Solutions to this problem, based on a method called "majority" or "quorum", are known, in which a certain number of license servers are configured to work cooperatively. As long as the majority of those servers is up and running and communicating with each other, all licenses are available, whereas as soon as the number of active license servers becomes less than the majority, all of the servers stop serving licenses. All existing solutions do not allow flexibility in the number of license servers

that participate in the cluster. This number is either fixed by the licensing system vendor or can be chosen upfront by the user when configuring the system, but it is not possible to increase or decrease the number of license servers in the cluster, during the life-cycle of the cluster itself. Having the possibility of adding and removing license servers to and from a cluster is an important feature to ensure the required flexibility for adapting the system capacity to the changing demands of the application environment.

In theory flexibility could be provided, for example, by binding the license authorization key to a software-based, random generated, binary identifier that can be securely stored into the license server's data base instead of binding them (the license key) to some specific license server hardware-based identifier. The same software-based binary identifier can be shared by all license servers participating in the cluster. However providing such a flexibility without any limitations on the way license servers can be added or removed from the cluster breaks the security of the license management system.

It is an object of the present invention to alleviate the above drawbacks of the prior art.

According to the present invention, we provide, in a network comprising a plurality of client workstations having a software program installed thereon, and an initial plurality of S license servers, a license management system for allowing the concurrent use of a maximum number n of copies of the software program, each client workstation requiring an authorization from one of the license servers for using the software program, the license management system requiring that at least the integer majority of the plurality of license servers is active at any time, the license management system comprising:

- means for allowing an increase or decrease in the number of license servers;

- means for limiting the number of the plurality of license servers with respect to the initial number S so that the number d of servers which can be added or removed from the initial number S respects the following rule:

$$2d < \text{integer majority of } S+d.$$

Furthermore, according to the present invention, we provide a method for providing flexibility to a license management system, the license management system permitting the concurrent use of n copies of a software program over a network comprising a plurality of client workstations, each client workstation having a copy of the software program installed thereon requiring an authorization from one of a plurality of S license servers each time the software program is used, the license management system requiring

that at least the integer majority of the plurality of license servers is active at any time, the method comprising the step of:

- allowing an increase or decrease in the number of license servers;
 - limiting the number of the plurality of license servers with respect to the initial number S so that the number d of license servers nodes which can be added or removed from the initial number S respects the following rule:
- $2d < \text{integer majority of } S+d.$

Also according to the present invention we provide a computer program product stored on a computer readable medium for allowing, in a network comprising a plurality of client workstations having a software program installed thereon, and an initial plurality of S license servers, the concurrent use of a maximum number n of copies of the software program, each client workstation requiring an authorization from one of the license servers before using the software program, the computer program product requiring that at least the integer majority of the plurality of license servers is active at any time, the computer program product comprising:

- computer readable program code means for allowing an increase or decrease in the number of license servers;
 - computer readable program code means for limiting the number of the plurality of license servers with respect to the initial number S so that the number d of servers which can be added or removed from the initial number S respects the following rule:
- $2d < \text{integer majority of } S+d.$

Various embodiments of the invention will now be described in detail by way of examples, with reference to accompanying figures, where:

Fig. 1 shows schematically an example of a license management system;

Fig. 2 and 3 are diagrams showing the functioning of a license management system;

Fig. 4 is an example of a license data base according to a preferred embodiment of the present invention;

Fig. 5 is an example of malicious use of the possibility of increasing the number of servers in the cluster.

As mentioned above Fig. 1 represents a typical network (e.g. a Local Area Network) using a license management system which could implement the present invention. Servers 101 may be, for example IBM RISC System/6000 43P-140 produced by International Business Machines Corporation. Client nodes 103 could be any personal computer or workstation available on the

market, e.g. IBM Personal Computer 300 GL produced by International Business Machines Corporation.

With reference to Fig. 2 and Fig. 3 the method of a license management system is described. Fig. 2 is a diagram of the method of functioning of a server according to a preferred embodiment of the present invention. The process starts at step 201 and goes straight to step 203 where the information about the software product to be licensed and the number of available licenses are read into the server memory. This information is usually provided by the software vendor and is usually protected against counterfeiting. When a request is received from a client (step 205) the server checks whether it is a request for a license (207). If this is the case and a license is available (209), the server creates a new license instance record (211) and decrements by one the number of available licenses. The server then sends a reply to the client (221) authorising the client to use the software product. Otherwise, if no more licenses are available the reply sent to the client (221) will be that the software product cannot be used. Going back to step 207, it may be the case that the client is requesting to release a license (215) after having used the software product. The server then deletes the correspondent license instance record (217) and increments the number of available licenses (219). With reference to Fig. 3 the functioning of a client wishing to use a software product is represented. When an available license server is found (303) a license request is issued to the server (305, 307). If the license is granted by the server then the client can use the software product (311), otherwise it is checked whether another server is available and control either goes back to step 305 or terminates the process. When the client finishes using the software application, a message is sent to the server which granted the license to release the license (313, 315).

Those skilled in the art will appreciate that a number of different method implementing similar license management procedures can be used instead of the one described above.

Fig. 4 shows an example of license data base 401, which, according to a preferred embodiment of the present invention, should be stored on each server 101. The information contained in this data base is:

- the cluster ID;
- the total number of servers belonging to the cluster;
- the secure ID (possibly hardware-based) of each license server that has ever participated to the cluster;
- the software products managed by the system;
- the number of available licenses for each product.

Those skilled in the art will appreciate that a method to lock the above described data base to the server hardware and to ensure the security of the

data base itself is needed. Furthermore those skilled in the art will appreciate that the above information could be organised in a number of different way according to well known programming techniques.

5 The network represented in Fig. 5a has three servers 101 serving a number n of licenses for software product 1 to a plurality of N clients 103. Each server 101 has a data base file 401 containing all the information described above. From the client point of view it does not make any difference whether a license is granted by server A, B or C. Any server 101 of the cluster can provide a license to any client 103. With such a system the failure of one of the servers would not be a problem, because the other two can do the service for all the clients. One of the servers could be inactive and act just as a backup server in case of failure of one of the other two. The limitation with such a configuration is that the majority of the servers (in this case two over three) must be always active, otherwise the system interrupts its services. Without this strict limitation the security of the system cannot be granted, because a malicious user could detach a server from the cluster, create another cluster and make believe to both cluster that they are still working in a three server cluster with, respectively one and two servers non-active. In this way the number of the possible licenses would be doubled. This protection mechanism is the "majority" requirement, well known by those skilled in the art.

25 A non limiting example of the system represented in Fig. 5a may be the following: a cluster of three servers serving 20 licences to 50 clients: the majority of the servers is two in this case, so at least two servers must be always active. Each of the server A and B manages 10 licenses each responding to the requests of the 50 clients, while server C stays non active. If one of the two active servers fails, then the server C takes over the licenses managed by that server and the service can continue without interruption and the majority rule is respected. Those skilled in the art will appreciate that a number of different implementations can be realised instead of the one described above, depending also from the requirements of the network.

35 As mentioned above, a desirable feature of a license management system is the flexibility of adding or deleting servers to the server cluster, in case circumstances change. In the system described in Fig. 5a the number of servers could be easily increased as shown in Fig. 5b. In this new configuration the data bases 401 in each server 101 have been updated to reflect the new situation. Each server "knows" that there are five servers belonging to the cluster and that at least three servers must always be contemporarily active to ensure the security of the system.

45 However a system which allowed a change of configuration as the one described with reference to Fig. 5a and 5b would be unsafe, and would not

guarantee the respect of the maximum number of licenses available with consequent damage for the software vendor who authorised the use of only n concurrent licenses. A malicious user could bypass the security check and act as illustrated in Fig. 5c. The malicious user could make a backup copy of the data base 401 of Fig. 5a when the cluster included three servers. Then the number of servers are increased to five to arrive in the configuration of Fig. 5b, already described. At this point the malicious user split the LAN 501 in two smaller ones not communicating with each other: one sub-LAN 505 including the three servers C, D and E; and a second sub-LAN 503 including the other two servers A and B. The sub-LAN 505 would continue its service because the Majority requirement is respected; each server believes that two servers are not active for some reason but, since the majority of servers is still available the service can be continued regularly by serving the n available licenses in the sub-LAN 505. The sub-LAN 503 would be interrupted, because the expected minimum number of three active server is not respected. However the malicious user restores the licensing environment, by substituting the data bases in sub-LAN 503 with the backup copies he made when the network was in the original configuration of Fig. 5a. At this point the two server A and B erroneously believe they are in a cluster of three servers with two of them (the Majority) active and they serve all the n licenses in the sub-LAN 503. In this way the malicious user has duplicated the number of available licenses causing a loss to the software vendor.

According to a preferred embodiment of the present invention this unwanted security exposure is avoided, while still allowing flexibility, by imposing a limitation in the number of servers that can be added or removed by the original configuration.

If we take S as the initial number of licence servers, d as the maximum number of servers that can be added or removed to or from the cluster and $M(x)$ as the integer majority of x , the limitation is the following: the difference between the maximum allowed number of servers belonging to the cluster ($S+d$) and the minimum allowed number of servers belonging to the cluster ($S-d$) must be less than the majority of the maximum number of server belonging to the cluster. In other words, the following relationship must be respected:

$$(S+d) - (S-d) < M(S+d), \text{ or} \\ 2d < M(S+d)$$

In the case of the example described above, where the initial number S of servers in the cluster was 3, the maximum number d of server that can be added or removed to the cluster respecting the rule above is 1.

According to a preferred embodiment of the present invention, once a licence server has been added to the cluster, its unique identification is

stored in the license data base of each server in the cluster and can never be deleted, unless the whole cluster and its identifier is deleted; this is to ensure that license servers cannot be replaced in the cluster. The license servers can only be added within the limits defined by the above rule and then deactivated to reduce the number of members of the cluster whose majority is to be up and running in order for the cluster to work. Once the limit specified above has been reached, non new server can be added to the cluster, but previously deactivated servers must be used. Allowing removal of license servers from the cluster or substitution would break the security of the cluster itself; for the same reason the unique ID of the license servers that initially form the cluster must be specified at cluster creation time.

The table below shows the values d of the maximum number of servers that can be added or deleted to and from the cluster for cluster having an initial number of servers between 4 and 10:

INITIAL CLUSTER SERVERS NUMBER S	MAX d
4	1
5	1
6	2
7	2
8	2
9	3
10	3

CLAIMS

1. In a network comprising a plurality of client workstations having a software program installed thereon, and an initial plurality of S license servers, a license management system for allowing the concurrent use of a maximum number n of copies of the software program, each client workstation requiring an authorization from one of the license servers for using the software program, the license management system requiring that at least the integer majority of the plurality of license servers is active at any time, the license management system comprising:

- means for allowing an increase or decrease in the number of license servers;

- means for limiting the number of the plurality of license servers with respect to the initial number S so that the number d of servers which can be added or removed from the initial number S respects the following rule:

$2d < \text{integer majority of } S+d.$

2. The license management system of claim 1 wherein each license server is able to authorise the use of a portion of said maximum number n of concurrent copies of software program.

3. The license management system of claim 1 or 2 further comprising:

- means for allocating to each server a unique ID, the maximum number of the IDs the system can allocate being S+d;

- means for locking the unique ID on each server;

- means for storing on each license server the ID of every other license server.

4. The licence management system of claim 3 further comprising:

- means for tracking license servers which are removed;

- means for preventing the addition of new servers once the maximum number of IDs that can be allocated has been reached.

5. A method for providing flexibility to a license management system, the license management system permitting the concurrent use of n copies of a software program over a network comprising a plurality of client workstations, each client workstation having a copy of the software program installed thereon requiring an authorization from one of a plurality of S license servers each time the software program is used, the license management system requiring that at least the integer majority of the plurality of license servers is active at any time, the method comprising the step of:

- allowing an increase or decrease in the number of license servers;

- limiting the number of the plurality of license servers with respect to the initial number S so that the number d of license servers which can be added or removed from the initial number S respects the following rule:
 $2d < \text{integer majority of } S+d.$

5

6. The method of claim 4 further comprising the steps of:
- allocating to each server a unique ID, the maximum number of the IDs the system can allocate being $S+d$;
- locking the unique ID on each server;
- storing on each license server the ID of every other license server.

10

7. The method of claim 6 wherein the step of tracking comprises:
- tracking license servers which are removed;
- preventing the addition of new servers once the maximum number of
IDs that can be allocated has been reached.

15

8. A computer program product stored on a computer readable medium for allowing, in a network comprising a plurality of client workstations having a software program installed thereon, and an initial plurality of S license servers, the concurrent use of a maximum number n of copies of the software program, each client workstation requiring an authorization from one of the license servers before using the software program, the computer program product requiring that at least the integer majority of the plurality of license servers is active at any time, the computer program product comprising:

20

25

- computer readable program code means for allowing an increase or decrease in the number of license servers;
- computer readable program code means for limiting the number of the plurality of license servers with respect to the initial number S so that the number d of servers which can be added or removed from the initial number S respects the following rule:
 $2d < \text{integer majority of } S+d.$

30

9. The computer program product of claim 8 further comprising:
- computer readable program code means for allocating to each server a unique ID, the maximum number of the IDs that can allocated being $S+d$;
- computer readable program code means for locking the unique ID on each server;
- computer readable program code means for storing on each license server the ID of every other license server.

35

40

10. The computer program product of claim 9 wherein the computer program code means for tracking comprises:
- computer program code means for tracking license servers which are removed;

45

- computer program code means for preventing the addition of new servers once the maximum number of IDs that can be allocated has been reached.



Application No: GB 9823187.1
Claims searched: 1-10

Examiner: K. Sylvan
Date of search: 24 March 1999

Patents Act 1977
Search Report under Section 17

Databases searched:

UK Patent Office collections, including GB, EP, WO & US patent specifications, in:

UK Cl (Ed.Q): G4A (AAP)

Int Cl (Ed.6): G06F (1/00)

Other: Online: EPODOC

Documents considered to be relevant:

Category	Identity of document and relevant passage	Relevant to claims
	NONE	

X	Document indicating lack of novelty or inventive step	A	Document indicating technological background and/or state of the art.
Y	Document indicating lack of inventive step if combined with one or more other documents of same category.	P	Document published on or after the declared priority date but before the filing date of this invention.
&	Member of the same patent family	E	Patent document published on or after, but with priority date earlier than, the filing date of this application.